The AbedGraham Group
Clinically Optimized Success

nccgroup

An AbedGraham and NCC Group Publication

The patient safety impact of network infrastructure vulnerabilities

Written by:

Dr Saif F Abed, Director of Cybersecurity Advisory Services, The AbedGraham Group
Dr Gabriel Ma, Senior Clinical Strategist, The AbedGraham Group
Stuart Kurutac, Senior Security Consultant, NCC Group

Contents

1 **Overview: IoT in healthcare**

2 **Methodology and use case**

3 **Results**

4 **Recommendations and conclusions**

5 **About The AbedGraham Group and NCC Group**

As a team of physicians and security analysts, The AbedGraham Group has been keeping tabs on the latest developments around IoT in healthcare. Our mix of clinical and technical expertise means that we are particularly well placed to judge cybersecurity events in the healthcare sector from a clinical safety and operational disruption perspective.

NCC Group is a global cyber security and risk mitigation expert that helps organisations across multiple sectors to assess, develop and manage cyber threats. It recognises that as industries become more digitised and connected, the complexity and interdependence of the threat landscape has never been higher.

As healthcare embraces digitisation, the volume of connected IoT devices contributing to the ecosystem is growing exponentially, and the level of interdependence that has been created is only increasing.

The AbedGraham Group and NCC Group have teamed up to provide our holistic view of this expanding medical IoT space, using our joint cyber security expertise to review the risks associated with connected devices in healthcare by looking at case studies, using quantitative analytics and providing recommendations for incident prevention.

IoT in healthcare

A greater attack surface area now exists as a result of the increasing number of connected devices present within healthcare providers' estates.

The presence of outdated operating systems across a large percentage of medical devices has provided an attractive target and presented more opportunities for attackers to try to breach the network and cause widescale disruption.

A majority of medical devices first became internet connected when vendors needed an efficient way to monitor and update their products remotely. Vendors eventually went on to improve the connectivity of these devices to other medical systems to overcome problems with poor interoperability and inefficiencies related to accessing and sharing siloed medical information. This increased connectivity has contributed to the increase in complexity and interdependency between medical systems, particularly increasing the risk of spread of contagion following a successful security breach.

This is also true when it comes to the connectivity of infrastructure-related assets, as enabling remote monitoring and management capabilities to these assets has contributed to the complexity of increased connectivity and interdependency within the healthcare ecosystem.

Risks of IoT in healthcare

Understandably, a greater attack surface area now exists as a result of the increasing number of connected devices present within healthcare providers' estates (including other infrastructure related assets, such as uninterruptable power supplies and building management systems etc). In addition, the presence of outdated operating systems across a large percentage of medical devices has provided an attractive target and presented more opportunities for attackers to try to breach the network and cause widescale disruption. Bad actors have a wide range of options when choosing which assets to compromise with varying results given their level of preparedness, the complexity of attack, and degree of disruption they wish to cause. Depending on the scenario, attacks could result in a minor delay to a patient awaiting care or cause an organisation to divert a large volume of patients to other external providers.

The hype surrounding IoT: clinical risk and concerns

With the increased level of medical device connectivity to mission-critical systems, it is reasonable for healthcare organisations to be concerned that compromise of these devices can have a more substantial and scalable impact, potentially affecting clinical delivery and the safety of patients at scale. This is also true for critical infrastructure assets that can be involved in supporting the delivery of clinical care at scale.

However, it is also important to recognise that as part of the wider holistic context, a majority of these medical devices will be used for individual patients at a time, can be easily swapped for another working device and will be accompanied by staff and under clinical supervision.

Obviously, key devices supporting mission-critical functions do exist, but the most significant risk for the larger proportion of connected medical devices will be as a point of entry for an attacker's reconnaissance purposes and to set the stage to help prepare for a larger scale attack at a later date.

A similar picture was observed following the discovery of the URGENT/11 in 2019, and then with Ripple20 earlier this year. Both sets of vulnerabilities affected different TCP/IP stacks used for many years in a diverse range of IoT devices, including connected medical devices. There were initial concerns that compromise of such a high volume of devices at scale could lead to significant patient harm, however it became apparent that an attack would have to be particularly targeted and complicated to compromise patient safety, making it highly unlikely to occur in live environments. Even then, the impact would likely be minimised due to attacks being involved in isolated cases and/or mitigated by clinical supervision.

Beyond these however, scalability in terms of clinical risk is less well documented when it comes to the role of network infrastructure. Accordingly, we decided to conduct a risk analysis of a critical vulnerability in this segment of the health system estate. The critical remote code execution vulnerability affecting several versions of F5's BIG-IP solutions was selected.

F5 BIG-IP

F5 BIG-IP is a collective name for a group of products offering a number of functions. These include traffic management and web application firewalls that provide availability and security to an organisation's applications and services. For example, a load balancer directing inbound traffic for a web application hosted across several servers. The load balancer can determine which server is able to facilitate additional requests based on the current volume each server is handling. Devices such as these are typically placed on an organisation's perimeter meaning they are Internet-facing, increasing the risk of attack.

The BIG-IP vulnerability that was disclosed on the 1st July 2020 was given a CVSS v2 score of 10, which is the highest score possible, and a CVSS v3 score of 9.8. The issue discovered allows unauthenticated users to read and write files, execute commands and disable services through the BIG IP Traffic Management User Interface (TMUI). This means if the TMUI is exposed to the Internet anyone that can access it would be able to take complete control of the device with full administrative privileges. The role that these devices provide could also facilitate further attacks and, depending on what the device is connected to, act as a pivot point into the internal network.

If the TMUI is exposed to the Internet anyone that can access it would be able to take complete control of the device with full administrative privileges.

The role that these devices provide could also facilitate further attacks and, depending on what the device is connected to, act as a pivot point into the internal network.

In a lot of cases the most difficult part of compromising an organisation's network is gaining the initial foothold.

Vulnerabilities such as this allow an attacker to gain that foothold in a privileged context without relying on interaction with anyone else.

Methodology and use case

In this whitepaper, The AbedGraham Group together with NCC Group has prepared a qualitative and quantitative analysis of the following use cases to provide a contextualised assessment, showcasing the cybersecurity risks that can be associated with IoT devices and network infrastructure in healthcare within non-clinical and clinical scenarios.

Use Case 1: What would be the risk associated with this highly critical vulnerability being present on a F5 BIG-IP Load Balancer which supports the handling and direction of inbound web traffic from its external facing websites for a healthcare system?

Use Case 2: What would be the risk associated with this highly critical vulnerability present on a F5 BIG-IP Load Balancer which supports domain authentication and bandwidth management through the authentication server providing access to the Electronic Health Record (EHR) clinical application within a healthcare system?

Use case analysis

We will detail the results of a holistic analysis of the associated cybersecurity risks starting with qualitative technical analysis of the risks by NCC Group, followed by a quantitative clinical assessment by The AbedGraham Group.

Qualitative technical review

In a lot of cases the most difficult part of compromising an organisation's network is gaining the initial foothold. A primary way to gain that foothold is through phishing, which requires interaction with an employee via email or some other, similar form of communication. However, vulnerabilities such as this allow an attacker to gain that foothold in a privileged context without relying on interaction with anyone else. The only caveat is that the interface of the device needs to be exposed to the Internet or reachable to the attacker by another equally trivial route.

The vulnerability is essentially a path traversal flaw that bypasses authentication of the TMUI and allows access to utility modules within the interface. These can be used to manipulate files, run BIG IP administrative commands or gain complete command execution on the device with full administrative privileges. The root cause of the issue stems from the differences in the way Apache httpd and Apache Tomcat handle the incoming request.

Consider the following URL:
https://[IP]/tmui/login.jsp/..;/tmui/locallb/workspace/fileRead.jsp?fileName=/etc/passwd

When issuing commands through a terminal, the use of "../" typically moves a user up one level in the directory hierarchy. For example, when using cmd.exe in Windows if a user is currently in "C:\Users\Bob\Documents" typing the change directory command (cd) followed by "../" at the command prompt will move that user to "C:\Users\Bob". This is also true for Unix based systems. Regarding the BIG- IP device, Apache httpd handles the initial request, because it allows paths containing a ";" without performing canonicalization it passes the unaltered request to Tomcat. The behaviour of Tomcat is different though, it removes the ";" and then normalises the path.

What happens is the original request is transformed to:
https://[IP]/tmui/login.jsp/../tmui/locallb/workspace/fileRead.jsp?fileName=/etc/passwd

and then subsequently to:
https://[IP]/tmui/tmui/locallb/workspace/fileRead.jsp?fileName=/etc/passwd

(A deeper dive into the analysis of the root cause can be found at https://research.nccgroup.com/2020/07/12/understanding-the-root-cause-of-f5-networks-k52145254-tmui-rce-vulnerability-cve-2020-5902/).

When we take in both use cases, from a technical perspective, the risk of compromise is high. Attacks leveraging this type of vulnerability will be widespread and largely indiscriminate in nature. The impact will vary depending on the end goal of the attack and the architecture of the network within which a vulnerable device is deployed.

At face value, a vulnerable device in scenario 1 could allow an attacker with control of the device to disrupt service to the externally facing web applications or redirect end users to an application under the attacker's control. In the latter instance if any of the web applications used a login, redirection to a malicious site would facilitate acquiring credentials. This is assuming servers under the attacker's control are reachable, the load balancer and the associated web servers are isolated in a separate network segment, access to the device is appropriately secured from the rest of the internal network and any firewall rules are sufficiently robust.

In scenario 2 however, the device aids with authentication to a critical system and can reach other resources on the internal network indicating the ability to pivot into the internal network. This could allow an attacker to move laterally throughout the network discovering and compromising vulnerable devices and systems greatly increasing the impact an attack would have.

Quantitative clinical analysis methodology

The quantitative assessment involves the use of The AbedGraham Group's clinical security analytics platform – [CCOM^2]. This platform contextually analyses, ranks and visualises each endpoint based on the risks they present to a health system clinically, organisationally, financially and in terms of regulatory compliance using a standardised 1-12 point scale. This is achieved using algorithmic models that take into account a broad range of behavioural attributes of network endpoints based on their functional behaviour across clinical workflows and associated interdependencies. In doing so, a granular asset profile can be determined and different types of attack can be modelled based on the detected vulnerabilities allowing the platform to determine the severity of any potential patient safety risks and their scalability.

The key thematic impact metrics are defined as follows:

- *Clinical Risk* pertains to the potential severity of patient harm that could occur
- *Organisational Risk* pertains to the level of clinical workflow disruption or service shut down that could occur
- *Financial Risk* pertains to the potential level of recovery and regulatory costs, as well as revenue losses that could occur
- *Regulatory Risk* pertains to the severity of intervention from regulators following disruption and degree of reputational damage

The patient safety and clinical workflow disruption risk metrics produced can be scaled to provide a total health system risk profile and the insights can ultimately guide any remediation strategies and application of security controls.

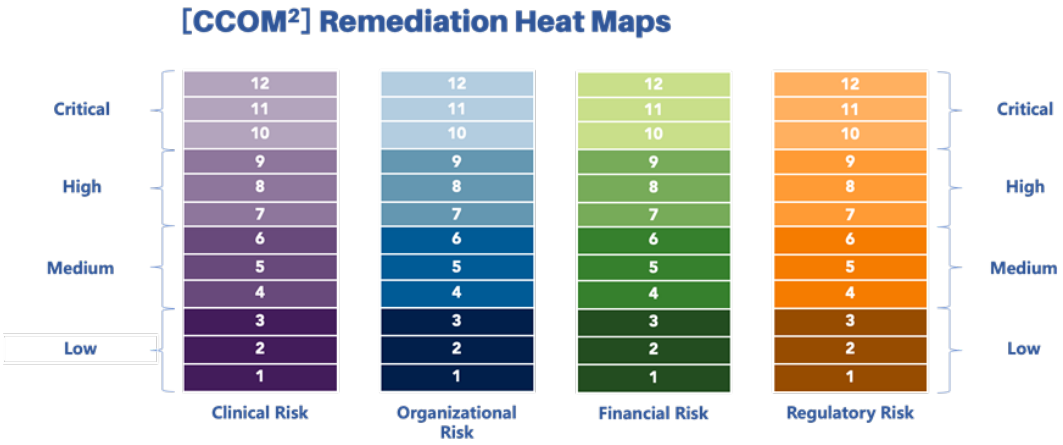

Visualisation of Risk Scores in [CCOM^2] with Heat Maps

Results - Use Case 1: F5 BIG-IP Load Balancer (inbound web traffic from external facing websites)

Quantitative results

Device	C	O	F	R	Total
Load Balancer	6	8	7	6	7

[CCOM^2] Risk Scores for a Load Balancer in Use Case 1, showing associated Clinical (C), Organisational (O), Financial (F), Regulatory (R), and Total Scores

Whilst a CVSS v3 score of 9.8 indicates that attackers may be able to easily exploit a vulnerability and cause harm in a number of ways remotely; in this scenario when contextualised within [CCOM^2], a load balancer with a vulnerability with a CVSS v3 score of 9.8 is associated with a high organisational risk score of 8/12. This is because failure of the device may directly affect administrative processes between community and hospital settings in terms of information sharing and appointment scheduling. As this endpoint is not involved in direct clinical care, but plays a role mainly as an interface between community and inpatient care, there may be knock on effects that would be associated with relatively moderate risk in the other clinical (6/12), financial (7/12) and regulatory (6/12) categories.

Discussion

By being connected to the network, there is always a baseline risk that any connected device can become a point of entry for spread into the wider network. Obviously, this risk can increase depending on a number of other factors (such as the type/level of connectivity involved and what other assets the device is connected to). In this case, the largest degree of disruption to the organisation would likely be to the administrative workflows and communication between the hospital and community should web traffic from the organisation's website be compromised. This may result in a delay in a certain cohort of patients making it into the hospital setting and receiving timely care. Moderate financial and regulatory risks would be likely due to delays in these patients being transferred to the hospital, potential re-scheduling of appointments and other associated revenue and remediation costs.

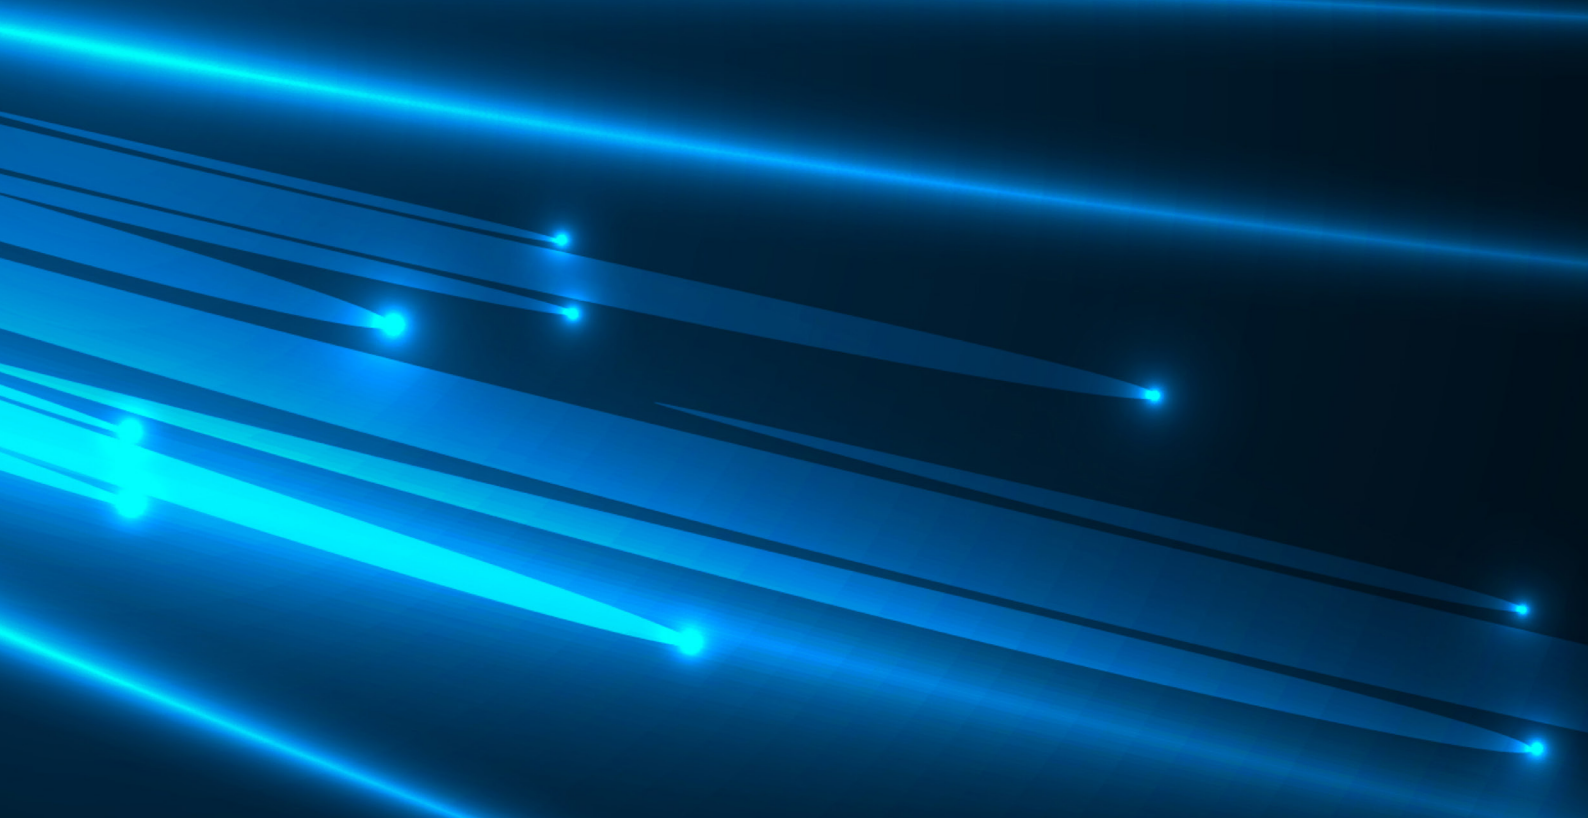

Results - Use Case 2: F5 BIG-IP Load Balancer (EHR authentication server)

Quantitative results

Device	C	O	F	R	Total
Clinical Load Balancer	9	9	8	8	8

[CCOM^2] Risk Scores for a Load Balancer in Use Case 2, showing associated Clinical (C), Organisational (O), Financial (F), Regulatory (R), and Total Scores

In this scenario when using [CCOM^2], a load balancer supporting clinical application uptime containing a vulnerability with a CVSS v3 score of 9.8, would understandably give a relatively high clinical risk and organisational risk score (both 9/12), as compromised accessibility to the mission-critical EHR would delay clinical information gathering and decision making (including formulation of management plans), impacting clinical workflows and affecting both staff and patients at scale across the organisation. Due to the decreased productivity across the entire estate (applicable to most clinical and administrative settings), this will have knock on effects that would be associated with relatively higher risks in the other financial (8/12) and regulatory (8/12) categories.

Discussion

The key issue associated with compromise of this asset with such a high scoring vulnerability is the resulting decrease in clinical productivity, resulting from a delay rather than cessation of service provision. As the asset is connected to a key mission-critical application that supports both clinical and administrative workflow at scale, the extent of disruption and risk of spread is of higher concern. Compared to Use Case 1, this would impact a wider range of activities across inpatient, outpatient and community related settings. Higher financial and regulatory risks would result from decreased productivity at scale across multiple environments, again including re-scheduling/re-organising clinical activities and other associated revenue and remediation costs.

Recommendations and conclusion

As healthcare organisations become increasingly digital and interconnected, it is vital that providers understand how security risks associated with network infrastructure truly impact clinical ecosystems. This means contextualising technical risks in a more nuanced way by taking into account clinical and organisational impacts, which in turn can also be translated into financial loss projections for healthcare providers.

As demonstrated by our case studies, although traditional concerns have been focused primarily on the risks associated with medical devices (as they can be involved in the provision of direct clinical care to patients), the impacts related to network infrastructure should not be overlooked.

In particular, these devices often have more far-reaching implications than first expected, particularly when they support mission-critical clinical applications and systems. With the volume and variety of endpoints spread across the healthcare estate only set to increase, the ability for healthcare providers to prepare their security teams to be able to quantify the associated risks, prioritise assets for remediation and decide how to allocate resources becomes even more crucial and unavoidable.

About NCC Group

NCC Group exists to make the world safer and more secure.

As global experts in cyber security and risk mitigation, NCC Group is trusted by over 15,000 clients worldwide to protect their most critical assets from the ever-changing threat landscape.

With the company's knowledge, experience and global footprint, it is best placed to help businesses identify, assess, mitigate and respond to the evolving cyber risks they face.

To support its mission, NCC Group continually invests in research and innovation, and is passionate about developing the next generation of cyber scientists.

About The AbedGraham Group

The AbedGraham Group is a global healthcare IT and cybersecurity technology group providing clinically led advisory services and analytics solutions for technology companies, government agencies and healthcare providers.

For further details about [CCOM^2] please visit and contact:

Website: www.abedgraham.com
E-mail: info@abedgraham.com
Twitter: @AbedGraham

www.nccgroup.com